Machine learning and stochastic control
Lecture 5

Jyväskylä Summer School
August 4-8, 2025

Huyên PHAM [1]

1. Ecole Polytechnique, CMAP

1. Fundamentals of SB and connection to optimal transport/control
2. SB as generative model

Introduction

• Schrödinger bridge (SB) problem historically addressed in 1931 as the dynamic transportation of cloud of particles.

• Formulated in modern (probabilistic) terms as the search of optimal process evolving from a measure to another one.

▶ SB has recently received great attention from the ML community as it is well-suitable for learning complex continuous time systems, notably for generative modeling

Schrödinger bridge (SB) problem

- $\Omega = C([0,T], \mathbb{R}^d)$ space of continuous $\mathbb{R}^d$-valued paths on $[0,T]$, X the canonical process, i.e. $X_t(\omega) = \omega(t)$, $t \in [0,T]$, $\mathcal{P}(\Omega)$ set of probability measures on path space Ω. For $\mathbb{P} \in \mathcal{P}(\Omega)$, $\mathbb{P}_t = X_t \# \mathbb{P} = \mathbb{P} \circ X_t^{-1}$: marginal law at t.
- μ_0 and μ_T two given probability measures on $\mathbb{R}^d$
- $\mathbb{Q}$ a prior/reference measure on Ω : belief of the dynamics before data observation, e.g., law of the Wiener process with initial measure ν_0.

(SBP) Find a measure $\mathbb{P}^*$ on path space solution to

$$\mathbb{P}^* \quad \in \quad \arg\min \left\{ \mathrm{KL}(\mathbb{P}|\mathbb{Q}) : \mathbb{P} \in \mathcal{P}(\Omega), \mathbb{P}_0 = \mu_0, \mathbb{P}_T = \mu_T \right\},$$

where

$$\mathrm{KL}(\mathbb{P}|\mathbb{Q}) \quad := \quad \begin{cases} \int \log\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right) \mathrm{d}\mathbb{P}, & \text{if } \mathbb{P} << \mathbb{Q} \\ \infty, & \text{otherwise} \end{cases}$$

is the Kullback-Leibler (or relative entropy) between two nonnegative measures.

Remark. By strict convexity of $\mathrm{KL}(\cdot|\mathbb{Q})$, the solution (when it exists) is unique.

Reduction to static SB

Denote by $\mathbb{P}_{0,T} = \mathbb{P} \circ (X_0, X_T)^{-1}$ the joint initial-terminal law of (X_0, X_T) under $\mathbb{P}$. We have the disintegration formula :

$$\mathbb{P}[\cdot] = \int \mathbb{P}^{xy}[\cdot]\mathbb{P}_{0,T}(\mathrm{d}x, \mathrm{d}y), \quad \text{with } \mathbb{P}^{xy}[\cdot] := \mathbb{P}[\cdot|(X_0, X_T) = (x, y)]$$

and similarly for $\mathbb{Q}$, $\mathbb{Q}^{xy}$, and $\mathbb{Q}_{0,T}$.

$\rightarrow$ Decomposition of the relative entropy :

$$\mathrm{KL}(\mathbb{P}|\mathbb{Q}) = \mathrm{KL}(\mathbb{P}_{0,T}|\mathbb{Q}_{0,T}) + \int\int \mathrm{KL}(\mathbb{P}^{xy}|\mathbb{Q}^{xy})\mathbb{P}_{0,T}(\mathrm{d}x, \mathrm{d}y)$$

- Reduction to static SB problem : minimize

$$\mathrm{KL}(\pi|\mathbb{Q}_{0,T}) = \int\int \log\left(\frac{\mathrm{d}\pi}{\mathrm{d}\mathbb{Q}_{0,T}}(x, y)\right)\pi(\mathrm{d}x, \mathrm{d}y) \tag{1}$$

over $\pi \in \Pi(\mu_0, \mu_T) = \{\pi \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d) : \pi_0 = \mu_0, \pi_T = \mu_T\}$.
The solution to **(SBP)** is then given by $\mathbb{P}^* = \int \mathbb{Q}^{xy}\pi^*(\mathrm{d}x, \mathrm{d}y)$ where π^* is solution to the static SBP.

Link with entropy-regularized optimal transport (OT)

- Consider $\mathbb{Q} = \mathbb{W}^\sigma$ the Wiener measure of variance σ^2, i.e., the law of the process

$$X_t = X_0 + \sigma W_t, \quad 0 \leq t \leq T, \quad X_0 \sim \nu_0,$$

with W a Brownian motion. In this case, $\mathbb{W}^\sigma_{0,T}(\mathrm{d}x, \mathrm{d}y) = \nu_0(\mathrm{d}x) q_\sigma(0, x, T, y)\mathrm{d}y$ where

$$q_\sigma(t, x, s, y) = \frac{1}{(2\pi\sigma^2(s-t))^{\frac{d}{2}}} \exp\Big(-\frac{|y-x|^2}{2\sigma^2(s-t)}\Big), \qquad 0 \leq t < s \leq T,\ x, y \in \mathbb{R}^d,$$

is the density transition kernel of $\mathbb{W}^\sigma$.

- For $\pi \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$, $\pi << \mathbb{W}^\sigma_{0,T}$, we denote by $\pi(x, y)$ its density w.r.t. $\nu_0(\mathrm{d}x)\mathrm{d}y$. Then,

$$\begin{aligned}\mathrm{KL}(\pi|\mathbb{W}^\sigma_{0,T}) &= \int\int \log\Big(\frac{\pi(x,y)}{q_\sigma(0,x,T,y)}\Big)\pi(\mathrm{d}x,\mathrm{d}y) \\ &= -\mathcal{E}(\pi) + \frac{1}{2\sigma^2 T}\int\int |x-y|^2 \pi(\mathrm{d}x,\mathrm{d}y) + \mathrm{cte},\end{aligned}$$

where $\mathcal{E}$ is the Shannon entropy :

$$\mathcal{E}(\pi) := -\int\int \log(\pi(x,y))\pi(\mathrm{d}x,\mathrm{d}y)$$

Entropic optimal transport (EOT)

Static SB problem is then equivalent to minimize

$$\begin{aligned} & \int\int |x-y|^2 \pi(\mathrm{d}x, \mathrm{d}y) - 2\sigma^2 T \mathcal{E}(\pi) \\ = \quad & \mathbb{E}_\pi |X_0 - X_T|^2 - \tau \mathcal{E}(\pi) \quad \text{over} \quad \pi \in \Pi(\mu_0, \mu_T) \end{aligned} \tag{2}$$

Problem (2) is an optimal transport problem with quadratic cost, regularized with the Shannon entropy via the parameter $\tau = \sigma^2 T$, and called entropic optimal transport.

$\rightarrow$ It allows to compute by **Sinkhorn algorithm** an approximation of the solution $\pi^{*,\tau}$ to EOT (2), which converges when $\tau \rightarrow 0$ to the solution π^* of classical OT :

$$\mathcal{W}_2^2(\mu_0, \mu_T) \quad := \quad \inf \left\{ \mathbb{E}_\pi |X_0 - X_T|^2 : \pi \in \Pi(\mu_0, \mu_T) \right\},$$

the (square) 2-Wasserstein distance between μ_0 and μ_T.

Diffusion process formulation of SB

Let $\mathbb{Q} = \mathbb{W}^\sigma$, and $\mathbb{P} \in \mathcal{P}(\Omega)$ such $\mathrm{KL}(\mathbb{P}|\mathbb{W}^\sigma) < \infty$. Then, there exists an $\mathbb{R}^d$-valued process α, adapted w.r.t. $\mathbb{F}$ the canonical filtration, with $\mathbb{E}_{\mathbb{P}}\big[\int_0^T \big|\frac{\alpha_t}{\sigma}\big|^2 \mathrm{d}t\big] < \infty$, s.t.

$$\begin{aligned} \frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{W}^\sigma} &= \frac{\mathrm{d}\mathbb{P}_0}{\mathrm{d}\nu_0} \exp\Big(\int_0^T \frac{\alpha_t}{\sigma}\frac{\mathrm{d}X_t}{\sigma} - \frac{1}{2}\int_0^T \big|\frac{\alpha_t}{\sigma}\big|^2 \mathrm{d}t\Big) \\ &= \frac{\mathrm{d}\mathbb{P}_0}{\mathrm{d}\nu_0} \exp\Big(\int_0^T \frac{\alpha_t}{\sigma}\mathrm{d}W_t^{\mathbb{P}} + \frac{1}{2}\int_0^T \big|\frac{\alpha_t}{\sigma}\big|^2 \mathrm{d}t\Big) \end{aligned} \tag{3}$$

and X follows the dynamics under $\mathbb{P}$ by Girsanov's theorem

$$\mathrm{d}X_t = \alpha_t \mathrm{d}t + \sigma \mathrm{d}W_t^{\mathbb{P}}, \quad 0 \le t \le T,$$

with $W^{\mathbb{P}}$ a Brownian motion under $\mathbb{P}$.

$\rightarrow$

$$\mathrm{KL}(\mathbb{P}|\mathbb{W}^\sigma) = \mathrm{KL}(\mathbb{P}_0|\nu_0) + \mathbb{E}_{\mathbb{P}}\Big[\frac{1}{2}\int_0^T \big|\frac{\alpha_t}{\sigma}\big|^2 \mathrm{d}t\Big]. \tag{4}$$

SB as a stochastic control problem

In view of (4), the **SBP** can be formulated equivalently as the stochastic control problem of finding the drift α of the diffusion

$$\mathrm{d}X_t \quad = \quad \alpha_t \mathrm{d}t + \sigma \mathrm{d}W_t,$$

that minimizes the energy functional

$$\mathbb{E}\Big[\frac{1}{2}\int_0^T \big|\frac{\alpha_t}{\sigma}\big|^2 \mathrm{d}t\Big]$$

under the marginal constraints : $X_0 \sim \mu_0$, $X_T \sim \mu_T$.

Case where $\mu_0 = \delta_{x_0}$, $\mu_T << $ Lebesgue measure λ

- We consider the prior measure $\mathbb{Q} = \mathbb{W}^\sigma$ with $\nu_0 = \delta_{x_0} = \mu_0$.

$\rightarrow$ The law of X_T under $\mathbb{W}^\sigma$ is $\nu_T(\mathrm{d}y) = q_\sigma(0, x_0, T, y)\mathrm{d}y$, i.e., the law of $\mathcal{N}(x_0, \sigma^2 T I_d)$

Theorem 1 (Solution to SB for Dirac initial law)

Assume that $\mathrm{KL}(\mu_T|\nu_T) < \infty$. The solution to SB is given by $\mathbb{P}^* = \psi(T, X_T)\mathbb{W}^\sigma$, where ψ satisfies the Kolmogorov equation :

$$\frac{\partial \psi}{\partial t} + \frac{\sigma^2}{2}\Delta\psi = 0, \quad t \in [0, T), \quad \text{and} \quad \psi(T, \cdot) = \frac{\mathrm{d}\mu_T}{\mathrm{d}\nu_T}. \tag{5}$$

where Δ is the Laplacian operator w.r.t. spatial variable. The corresponding optimal drift of the diffusion process X under $\mathbb{P}^*$ is

$$\alpha_t^* = \sigma^2 \nabla_x \log \psi(t, X_t), \quad 0 \leq t \leq T,$$

and the optimal cost is $\mathrm{KL}(\mathbb{P}^*|\mathbb{W}^\sigma) = \mathrm{KL}(\mu_T|\nu_T)$.

Remark. The PDE (5) means that $\psi(t, X_t)$ is a martingale under $\mathbb{W}^\sigma$, given by Feynman-Kac formula :

$$\psi(t, X_t) = \mathbb{E}_{\mathbb{W}^\sigma}\Big[\frac{\mu_T(X_T)}{q_\sigma(0, x_0, T, X_T)}\Big|\mathcal{F}_t\Big], \quad 0 \leq t \leq T. \tag{6}$$

Proof of Theorem 1 : law constraint under $\mathbb{P}^*$

- Let $\mathbb{P}^* = \psi(T, X_T)\mathbb{W}^\sigma$ with ψ given by (5). Then, for any bounded function f on $\mathbb{R}^d$,

$$
\begin{aligned}
\mathbb{E}_{\mathbb{P}^*}\big[f(X_T)\big] &= \mathbb{E}_{\mathbb{W}^\sigma}\Big[\frac{\mu_T(X_T)}{q_\sigma(0, x_0, T, X_T)} f(X_T)\Big] \\
&= \int \frac{\mu_T(y)}{q_\sigma(0, x_0, T, y)} f(y) q_\sigma(0, x_0, T, y)\mathrm{d}y \\
&= \int f(y)\mu_T(y)\mathrm{d}y,
\end{aligned}
$$

where we use in the second equality the fact that $y \mapsto q_\sigma(0, x_0, T, y)$ is the density of X_T under $\mathbb{W}^\sigma$. This proves that the law of X_T under $\mathbb{P}^*$ is μ_T.

- Moreover, by definition of the density $\psi(T, X_T)$ of $\mathbb{P}^*$, we have

$$
\begin{aligned}
\mathrm{KL}(\mathbb{P}^*, \mathbb{W}^\sigma) &= \mathbb{E}_{\mathbb{P}^*}\Big[\log\Big(\frac{\mu_T(X_T)}{q_\sigma(0, x_0, T, X_T)}\Big)\Big] \\
&= \int \log\Big(\frac{\mu_T(y)}{q_\sigma(0, x_0, T, y)}\Big)\mu_T(\mathrm{d}y) \\
&= \mathrm{KL}\big(\mu_T | \mathcal{N}(x_0, \sigma^2 T I_d)\big).
\end{aligned}
$$

Proof of Theorem 1 : drift under $\mathbb{P}^*$

- By Itô's formula, the positive martingale $Z_t := \psi(t, X_t)$ under $\mathbb{W}^\sigma$ satisfies

$$\mathrm{d}Z_t \quad = \quad \nabla_x \psi(t, X_t)\mathrm{d}X_t \; = \; Z_t \nabla_x \log \psi(t, X_t)\mathrm{d}X_t,$$

hence is in the Doléans-Dade exponential form (3) with

$$\alpha_t^* \quad = \quad \sigma^2 \nabla_x \log \psi(t, X_t).$$

- Moreover, by (4), we have

$$\mathrm{KL}(\mathbb{P}^* | \mathbb{W}^\sigma) \quad = \quad \mathbb{E}_{\mathbb{P}^*}\Big[\frac{1}{2}\int_0^T \big|\frac{\alpha_t^*}{\sigma}\big|^2 \mathrm{d}t\Big].$$

Proof of Theorem 1 : optimality of $\mathbb{P}^*$

- Let $\mathbb{P} \in \mathcal{P}(\Omega)$ associated with drift α, and s.t. $X_T \sim \mu_T$ under $\mathbb{P}$. Then, by (3)

$$
\begin{aligned}
1 &= \mathbb{E}_{\mathbb{W}^\sigma}\big[\psi(T, X_T)\big] \\
&= \mathbb{E}_{\mathbb{P}}\Big[\psi(T, X_T)\exp\Big(-\int_0^T \frac{\alpha_t}{\sigma}\mathrm{d}W_t^{\mathbb{P}} - \frac{1}{2}\int_0^T \Big|\frac{\alpha_t}{\sigma}\Big|^2 \mathrm{d}t\Big)\Big] \\
&\geq \exp\Big(\mathbb{E}_{\mathbb{P}}\Big[\log\psi(T, X_T) - \int_0^T \frac{\alpha_t}{\sigma}\mathrm{d}W_t^{\mathbb{P}} - \frac{1}{2}\int_0^T \Big|\frac{\alpha_t}{\sigma}\Big|^2 \mathrm{d}t\Big]\Big) \qquad \text{(by Jensen)} \\
&= \exp\Big(\mathrm{KL}\big(\mu_T | \mathcal{N}(x_0, \sigma^2 T)\big) - \mathbb{E}_{\mathbb{P}}\Big[\frac{1}{2}\int_0^T \Big|\frac{\alpha_t}{\sigma}\Big|^2 \mathrm{d}t\Big]\Big),
\end{aligned}
$$

which shows that

$$
\mathrm{KL}(\mathbb{P}|\mathbb{W}^\sigma) = \mathbb{E}_{\mathbb{P}}\Big[\frac{1}{2}\int_0^T \Big|\frac{\alpha_t}{\sigma}\Big|^2 \mathrm{d}t\Big] \geq \mathrm{KL}\big(\mu_T | \mathcal{N}(x_0, \sigma^2 T I_d)\big) = \mathrm{KL}(\mathbb{P}^*|\mathbb{W}^\sigma).
$$

□

Case of general distribution μ_0, $\mu_T << \lambda$

- Consider a prior measure $\mathbb{W}^\sigma$ with $\nu_0 = \lambda$, called *reversible Brownian motion* :
  - $\mathbb{W}^\sigma \circ X_t^{-1} = \lambda$.
  - The law (density) of $X_t|X_s = y$, for $t < s$, is $x \mapsto q_\sigma(t, x, s, y)$.

Theorem 2 (solution to SBP)

The solution to **SB** is given by $\mathbb{P}^* = \varphi(0, X_0)\psi(T, X_T)\mathbb{W}^\sigma$, where φ, ψ are non-negative functions satisfying the Schrödinger system :

$$\varphi(0,x)\psi(0,x) \;=\; \frac{\mathrm{d}\mu_0}{\mathrm{d}\lambda}, \;\; \lambda - \text{a.e.}, \qquad \varphi(T,y)\psi(T,y) \;=\; \frac{\mathrm{d}\mu_T}{\mathrm{d}\lambda}, \;\; \lambda - \text{a.e.},$$

with the Fokker-Planck, and Kolmogorov equations :

$$\frac{\partial \varphi}{\partial s}(s,y) - \frac{\sigma^2}{2}\Delta\varphi(s,y) \;=\; 0, \qquad \frac{\partial \psi}{\partial t}(t,x) + \frac{\sigma^2}{2}\Delta\psi(t,x) \;=\; 0. \tag{7}$$

The corresponding optimal drift of the diffusion process X under $\mathbb{P}^*$ is

$$\alpha_t^* \;=\; \sigma^2 \nabla_x \log \psi(t, X_t), \quad 0 \le t \le T.$$

About Schrödinger system

• The probabilistic representation of the Fokker-Planck and Kolmogorov equations (7) (Feynman-Kac formulae) are

$$\begin{aligned}\varphi(s,y) &= \mathbb{E}_{\mathbb{W}^\sigma}\left[\varphi(0,X_0)|X_s=y\right] = \int q_\sigma(0,x,s,y)\varphi(0,x)\mathrm{d}x\\ \psi(t,x) &= \mathbb{E}_{\mathbb{W}^\sigma}\left[\psi(T,X_T)|X_t=x\right] = \int q_\sigma(t,x,T,y)\psi(T,y)\mathrm{d}y,\end{aligned}$$

$\rightarrow$ Schrödinger system is written equivalently as the search of functions $\varphi^* = \varphi(0,\cdot), \psi^* = \psi(T,\cdot)$ s.t.

$$\begin{cases}\varphi^*(x)\mathbb{E}_{\mathbb{W}^\sigma}\left[\psi^*(X_T)|X_0=x\right] = \frac{\mathrm{d}\mu_0}{\mathrm{d}\lambda}, & \lambda-\text{a.e.},\\ \psi^*(y)\mathbb{E}_{\mathbb{W}^\sigma}\left[\varphi^*(X_0)|X_T=y\right] = \frac{\mathrm{d}\mu_T}{\mathrm{d}\lambda}, & \lambda-\text{a.e.}.\end{cases} \tag{8}$$

The solution to SB problem is then given by $\mathbb{P}^* = \varphi^*(X_0)\psi^*(X_T)\mathbb{W}^\sigma$.

• The functions φ^*, ψ^* represent the dual functions associated to the initial and terminal marginal law constraints of the convex entropic minimization problem. Existence is proved from the dual problem, and numerical approximation is achieved by Fortet-Sinkhorn algorithm. There is uniqueness up to a multiplication of φ^* and division of ψ^* by the same constant.

Proof of Theorem 2 : dual approach on static SB

Consider the Lagrangian function of the static **SB** problem (1) :

$$\begin{aligned}\mathcal{L}(\pi, \ell, \kappa) &= \int\int \big[\log\big(\frac{\pi(x,y)}{q_\sigma(0,x,T,y)}\big)\big]\pi(x,y)\mathrm{d}x\mathrm{d}y \\ &\quad + \int \ell(x)\big[\int \pi(x,y)\mathrm{d}y - \mu_0(x)\big]\mathrm{d}x + \int \kappa(y)\big[\int \pi(x,y)\mathrm{d}x - \mu_T(y)\big]\mathrm{d}y.\end{aligned}$$

The first order condition (pointwise derivative w.r.t. $\pi(x,y)$) yields :

$$1 + \log\big(\frac{\pi^*(x,y)}{q_\sigma(0,x,T,y)}\big) + \ell(x) + \kappa(y) = 0, \quad \lambda \otimes \lambda - \text{a.e.}$$

$\rightarrow$

$$\frac{\pi^*(x,y)}{q_\sigma(0,x,T,y)} = \varphi^*(x)\psi^*(y)$$

for some functions $\varphi^*(x) = e^{-1-\ell(x)}$ of x, and $\psi^*(y) = e^{-\kappa(y)}$ of y.

Proof of Theorem 2 : marginal law constraints

By writing that $\pi^* \in \Pi(\mu_0, \mu_T) \to \varphi^*, \psi^*$ have to satisfy the Schrödinger system :

$$\begin{aligned} \varphi^*(x) \int q_\sigma(0, x, T, y)\psi^*(y)\mathrm{d}y &= \mu_0(x), \qquad \lambda - \text{a.e.}, \\ \psi^*(y) \int q_\sigma(0, x, T, y)\varphi^*(x)\mathrm{d}x &= \mu_T(y), \qquad \lambda - \text{a.e.} \end{aligned}$$

- The associated minimal relative entropy is

$$\begin{aligned} \mathrm{KL}(\pi^* | \mathbb{W}^\sigma_{0T}) &= \int\int \Big[\log\Big(\frac{\pi^*(x, y)}{q_\sigma(0, x, T, y)}\Big)\Big]\pi^*(x, y)\mathrm{d}x\mathrm{d}y \\ &= \int \big[\log \varphi^*(x)\big]\mu_0(x)\mathrm{d}x + \int \big[\log \psi^*(y)\big]\mu_T(y)\mathrm{d}y. \end{aligned}$$

Proof of Theorem 2 : back to the dynamic SB

• The solution to **SBP** is then : $\mathbb{P}^* = \varphi^*(X_0)\psi^*(X_T)\mathbb{W}^\sigma$, with associated relative entropy :

$$\mathrm{KL}(\mathbb{P}^*|\mathbb{W}^\sigma) = \mathbb{E}_{\mu_0}\big[\log\varphi^*(X_0)\big] + \mathbb{E}_{\mu_T}\big[\log\psi^*(X_T)\big].$$

• The optimal drift of X under $\mathbb{P}^*$ is

$$\alpha_t^* = \sigma^2\nabla_x \log\psi(t, X_t), \quad 0 \leq t \leq T,$$

with

$$\begin{aligned}\psi(t,x) &= \mathbb{E}_{\mathbb{W}^\sigma}\big[\psi^*(X_T)|X_t = x\big] \\ &= \int q_\sigma(t,x,T,y)\psi^*(y)\mathrm{d}y.\end{aligned}$$

□

Numerical approximation of the Schrödinger system

• Set $q(x, y) := q_\sigma(0, x, T, y)$ the reference (density) measure, which is strictly positive and bounded from above. The Schrödinger system is written as :

$$\begin{cases} \varphi^*(x) \int q(x, y)\psi^*(y)\mathrm{d}y & = \mu_0(x), \\ \psi^*(y) \int q(x, y)\varphi^*(x)\mathrm{d}x & = \mu_T(y). \end{cases} \tag{9}$$

▶ The continuous time version of the Sinkhorn (discrete space case) algorithm would consist in solving (9) by iteration with the following updates :

$$\varphi^{n+1}(x) = \frac{\mu_0(x)}{\int q(x,y)\psi^n(y)\mathrm{d}y}, \qquad \psi^{n+1}(y) = \frac{\mu_T(y)}{\int q(x,y)\varphi^{n+1}(x)\mathrm{d}x}, \tag{10}$$

initialized with some positive function ψ^0, e.g., $\psi^0 = 1_d$.

Remark. The map Φ s.t. $\varphi^{n+1} = \Phi[\varphi^n]$ in the scheme (10) does not have the good suitable properties for proving the fixed point property of Φ in the continuous time case. Fortet proposed a truncated modified scheme for getting the convergence of the scheme.

Iterative Proportional Fitting

Define the (density) probability mesures π^m, $m \geq 1$, by

$$\begin{aligned} \pi^{2n+1}(x,y) &:= \varphi^{n+1}(x)q(x,y)\psi^n(y), \\ \pi^{2n+2}(x,y) &:= \varphi^{n+1}(x)q(x,y)\psi^{n+1}(y), \end{aligned}$$

starting from some π^0, e.g., $\pi^0 = q$. Then, from dual approach, se see that

$$\begin{aligned} \pi^{2n+1} &\in \arg\min\left\{\mathrm{KL}(\pi|\pi^{2n}) : \pi_0 = \mu_0\right\}, \\ \pi^{2n+2} &\in \arg\min\left\{\mathrm{KL}(\pi|\pi^{2n+1}) : \pi_T = \mu_T\right\}. \end{aligned}$$

This iteration over π via the minimization of relative entropy where one alternately updates either with the initial or terminal law constraint, is called *Iterative Proportional Fitting* (IPF) scheme.

Generative AI

- We observe i.i.d. data from an unknown distribution $\mu_{data} \in \mathcal{P}(\mathbb{R}^d)$
- Goal : generate/synthesize new (real-looking) samples of μ_{data}

Figure – Real-looking images generated by AI : DALL-E, Midjourney, Stable diffusion, etc

Generative learning

Generative modeling (GM) has become a classical task in machine learning with several competing methods :

- **Likelihood-based models** : energy-based models (EBM), variational auto-encoders (VAE)
- **Implicit generative models** : generative adversarial network (GAN)
- **Diffusion models** :
  - *Score-based models* (2020-) : (from the blog of Y. Song)

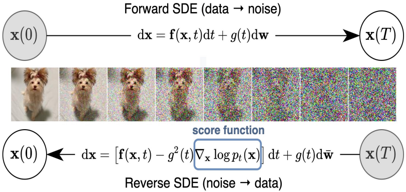

  - *Optimal transport/SB approach* (2022-) : this lecture

Generative learning via Schrödinger bridge (SB)

- De Bortoli, Thornton, Heng, Doucet (21) propose a noising step from SB with $\mu_0 = \mu_{data}$ data distribution, $\mu_T = \mu_{noise}$ prior distribution generated over **finite time**
- Wang et al. (22) : learn $\mu_{data} << \lambda$ in two stages

1. Learn a smooth version of μ_{data} (noising with Gaussian distribution) :

$$\mu_\varepsilon(x) = \int \mu_{data}(y)\phi_\varepsilon(x-y)\mathrm{d}y,$$

   where ϕ_ε is the density of $\Phi_\varepsilon \equiv \mathcal{N}(0, \varepsilon^2 I_d)$.

2. Learn μ_{data} starting from μ_ε.

Both steps are performed via Schrödinger bridge diffusion on $[0,1]$ ($T = 1$).

Analytic two-stage solution to the SB

Theorem 3

1. Define $\psi_1(x) = \frac{\mathrm{d}\mu_\varepsilon}{\mathrm{d}\Phi_\sigma}(x)$. Then, for the SDE

$$\mathrm{d}X_t = \sigma^2 \nabla_x \log \mathbb{E}_{\mathbb{W}^\sigma}\left[\psi_1(X_1)|X_t\right]\mathrm{d}t + \sigma \mathrm{d}W_t, \quad 0 \leq t \leq 1, \tag{11}$$

   starting from $X_0 = 0$, we have $X_1 \sim \mu_\varepsilon$.

2. For the SDE

$$\mathrm{d}X_t = \varepsilon^2 \nabla_x \log \mu_{\sqrt{1-t}\varepsilon}(X_t)\mathrm{d}t + \varepsilon \mathrm{d}W_t, \quad 0 \leq t \leq 1, \tag{12}$$

   starting from $X_0 \sim \mu_\varepsilon$, we have $X_1 \sim \mu_{data}$.

Comment : The target distribution μ_{data} can be learned from Dirac Mass through two SDEs (11)-(12) on finite time interval, that are simulated by Euler scheme once we are able to estimate the drift terms.

Proof of Theorem 3

- SDE (11) follows directly from Theorem 1.

- For SDE (12), we solve the Schrödinger system in Theorem 2 : we see that $\varphi^*(x) \equiv 1$, $\psi^*(y) = \mu_{data}(y)$ solve the Schrödinger system (8) (with diffusion coefficient ε) since

$$\begin{cases} \varphi^*(x)\mathbb{E}_{\mathbb{W}^\varepsilon}\left[\psi^*(X_1)|X_0 = x\right] = \mathbb{E}_{\mathbb{W}^\varepsilon}\left[\mu_{data}(X_1)|X_0 = x\right] = \mu_\varepsilon(x) \\ \psi^*(y)\mathbb{E}_{\mathbb{W}^\varepsilon}\left[\varphi^*(X_0)|X_T = y\right] = \mu_{data}(y). \end{cases}$$

The drift of the SDE (12), is then given by $\alpha_t^{2,*} = \varepsilon^2 \nabla_x \log \psi(t, X_t)$ with

$$\psi(t,x) = \mathbb{E}_{\mathbb{W}^\varepsilon}\left[\psi^*(X_1)|X_t = x\right] = \mathbb{E}_{\mathbb{W}^\varepsilon}\left[\mu_{data}(X_1)|X_t = x\right] = \mu_{\sqrt{1-t}\varepsilon}(x).$$

□

Estimation of the drift terms from data samples (1)

The drift term of SDE (11) is $\alpha_t^{1,*} = \sigma^2 S_1(t, X_t)$ where $S_1(t,x) := \nabla_x \log \psi_1(t,x)$ with

$$\begin{aligned}
\psi_1(t,x) &= \mathbb{E}_{\mathbb{W}^\sigma}\left[\psi_1(X_1)|X_t = x\right] = \mathbb{E}_{\mathbb{W}^\sigma}\left[\frac{\mu_\varepsilon}{\phi_\sigma}(X_1)|X_t = x\right] \\
&= \int \frac{\mu_\varepsilon}{\phi_\sigma}(y) q_\sigma(t,x,1,y)\mathrm{d}y = \int \frac{q_\sigma(t,x,1,y)}{\phi_\sigma(y)}\mu_\varepsilon(y)\mathrm{d}y \\
&= \mathbb{E}_{X_1 \sim \mu_{data}, Z \sim \mathcal{N}(0, I_d)}\left[\frac{q_\sigma(t,x,1,X_1+\varepsilon Z)}{\phi_\sigma(X_1+\varepsilon Z)}\right],
\end{aligned}$$

(q_σ is the transition density of σW, i.e., $q_\sigma(t,x,1,y) = \frac{1}{(2\pi\sigma^2(1-t))^{\frac{d}{2}}} \exp(-\frac{|y-x|^2}{2\sigma^2(1-t)})$)
$\rightarrow$

$$S_1(t,x) = \frac{\nabla_x \psi_1(t,x)}{\psi_1(t,x)} = \frac{\mathbb{E}_{X_1 \sim \mu_{data}, Z \sim \mathcal{N}(0, I_d)}\left[\frac{\nabla_x q_\sigma(t,x,1,X_1+\varepsilon Z)}{\phi_\sigma(X_1+\varepsilon Z)}\right]}{\mathbb{E}_{X_1 \sim \mu_{data}, Z \sim \mathcal{N}(0, I_d)}\left[\frac{q_\sigma(t,x,1,X_1+\varepsilon Z)}{\phi_\sigma(X_1+\varepsilon Z)}\right]}.$$

- The score function S_1 can be estimated based on observation samples X_1 of μ_{data}, and samples of $Z \sim \mathcal{N}(0, I_d)$.

Estimation of the drift terms from data samples (2)

Estimation of the time varying drift term of SDE (12), i.e., $\varepsilon^2 \nabla \log \mu_{\sqrt{1-t\varepsilon}}(x)$.

- Build a DNN $\mathrm{s}_\theta(x, \tilde{\varepsilon})$ that approximates the score $\nabla \log \mu_{\tilde{\varepsilon}}(x)$ for $\tilde{\varepsilon}$ varying in $[0, \varepsilon]$.
  - Score matching by minimizing

$$J_{SM}(\theta) \quad := \quad \mathbb{E}_{\mu_{\tilde{\varepsilon}}} \left| \mathrm{s}_\theta(\tilde{X}, \tilde{\varepsilon}) - \nabla \log \mu_{\tilde{\varepsilon}}(\tilde{X}) \right|^2.$$

  - Denoising score matching (Vincent 11) : this is equivalent to minimize Annex

$$J_{DSM}(\theta) \quad := \quad \mathbb{E}_{X \sim \mu_{data}, Z \sim \mathcal{N}(\mathbf{0}, \tilde{\varepsilon}^2 I_d)} \left| \mathrm{s}_\theta(X + Z, \tilde{\varepsilon}) - \nabla \log \phi_{\tilde{\varepsilon}}(Z) \right|^2,$$

where $\nabla \log \phi_{\tilde{\varepsilon}}(z) = -\frac{z}{\tilde{\varepsilon}^2}$ is the score of $\mathcal{N}(0, \tilde{\varepsilon}^2 I_d)$.

$\rightarrow$ Estimate s_θ by minimizing

$$\frac{1}{MN} \sum_{j=1}^{M} \sum_{i=1}^{N} \tilde{\varepsilon}_j^2 \left| \mathrm{s}_\theta(X_i + Z_{ij}, \tilde{\varepsilon}_j) + \frac{Z_{ij}}{\tilde{\varepsilon}_j^2} \right|^2$$

from i.i.d. samples $\tilde{\varepsilon}_j$, $j = 1, \ldots, M$, of $\mathcal{U}([0, \varepsilon])$, X_i, $i = 1, \ldots, N$, of μ_{data}, and Z_{ij}, from $\mathcal{N}(0, \tilde{\varepsilon}_j^2 I_d)$.

Numerical experiments

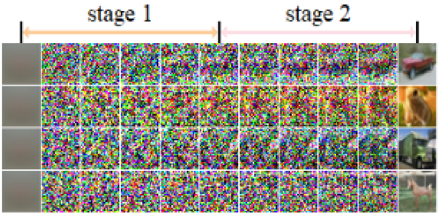

Figure – Particle evolution after two stages

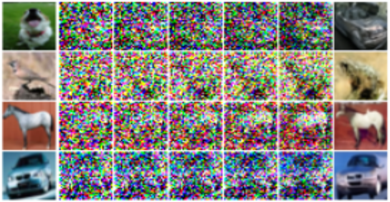

Figure – Denoising with stage 2 from perturbed real images

Annex : Proof of Denoising score matching | back

Fix $\tilde{\varepsilon}$, and set $\mathrm{s}_\theta(x) = \mathrm{s}_\theta(x, \tilde{\varepsilon})$. Recall $\mu_{\tilde{\varepsilon}}(\tilde{x}) = \int \mu_{data}(x)\phi_{\tilde{\varepsilon}}(\tilde{x} - x)\mathrm{d}x$.

- We develop the score matching criterion as

$$J_{SM}(\theta) \quad := \quad \mathbb{E}_{\mu_{\tilde{\varepsilon}}} \left| \mathrm{s}_\theta(\tilde{X}) - \nabla \log \mu_{\tilde{\varepsilon}}(\tilde{X}) \right|^2 = \mathbb{E}_{\mu_{\tilde{\varepsilon}}} \left| \mathrm{s}_\theta(\tilde{X}) \right|^2 - 2S(\theta) + C_1, \tag{13}$$

where $C_1 := \mathbb{E}_{\mu_{\tilde{\varepsilon}}} \left| \nabla \log \mu_{\tilde{\varepsilon}}(\tilde{X}) \right|^2$, and

$$\begin{aligned}
S(\theta) &= \mathbb{E}_{\mu_{\tilde{\varepsilon}}} \left[< \mathrm{s}_\theta(\tilde{X}), \nabla \log \mu_{\tilde{\varepsilon}}(\tilde{X}) > \right] = \int < \mathrm{s}_\theta(\tilde{x}), \nabla \log \mu_{\tilde{\varepsilon}}(\tilde{x}) > \mu_{\tilde{\varepsilon}}(\tilde{x})\mathrm{d}\tilde{x} \\
&= \int < \mathrm{s}_\theta(\tilde{x}), \nabla \mu_{\tilde{\varepsilon}}(\tilde{x}) > \mathrm{d}\tilde{x} = \int < \mathrm{s}_\theta(\tilde{x}), \frac{\partial}{\partial \tilde{x}} \int \mu_{data}(x)\phi_{\tilde{\varepsilon}}(\tilde{x} - x)\mathrm{d}x > \mathrm{d}\tilde{x} \\
&= \int < \mathrm{s}_\theta(\tilde{x}), \int \mu_{data}(x) \frac{\partial}{\partial \tilde{x}} \phi_{\tilde{\varepsilon}}(\tilde{x} - x)\mathrm{d}x > \mathrm{d}\tilde{x} \\
&= \int < \mathrm{s}_\theta(\tilde{x}), \int \mu_{data}(x)\phi_{\tilde{\varepsilon}}(\tilde{x} - x) \frac{\partial}{\partial \tilde{x}} \log \phi_{\tilde{\varepsilon}}(\tilde{x} - x)\mathrm{d}x > \mathrm{d}\tilde{x} \\
&= \int \int \mu_{data}(x)\phi_{\tilde{\varepsilon}}(\tilde{x} - x) < \mathrm{s}_\theta(\tilde{x}), \frac{\partial}{\partial \tilde{x}} \log \phi_{\tilde{\varepsilon}}(\tilde{x} - x) > \mathrm{d}x\mathrm{d}\tilde{x} \\
&= \mathbb{E}_{X \sim \mu_{data}, Z \sim \mathcal{N}(\mathbf{0}, \tilde{\varepsilon}^2 I_d)} \left[< \mathrm{s}_\theta(X + Z), \nabla \log \phi_{\tilde{\varepsilon}}(Z) > \right].
\end{aligned} \tag{14}$$

Annex : Proof of Denoising score matching (end) back

- On the other hand, the denoising score matching criterion is developed as

$$\begin{aligned} J_{DSM}(\theta) &:= \mathbb{E}_{X\sim\mu_{data}, Z\sim\mathcal{N}(\mathbf{0},\tilde{\varepsilon}^2 I_d)} \left| \mathrm{s}_\theta(X+Z) - \nabla \log \phi_{\tilde{\varepsilon}}(Z) \right|^2 \\ &= \mathbb{E}_{\mu_{\tilde{\varepsilon}}} \left| \mathrm{s}_\theta(\tilde{X}) \right|^2 - 2\mathbb{E}_{X\sim\mu_{data}, Z\sim\mathcal{N}(\mathbf{0},\tilde{\varepsilon}^2 I_d)} \left[< \mathrm{s}_\theta(X+Z), \nabla \log \phi_{\tilde{\varepsilon}}(Z) > \right] + C_2, \end{aligned} \tag{15}$$

where $C_2 := \mathbb{E}_{X\sim\mu_{data}, Z\sim\mathcal{N}(\mathbf{0},\tilde{\varepsilon}^2 I_d)} \left| \nabla \log \phi_{\tilde{\varepsilon}}(Z) \right|^2$.

- It follows from (13), (14) and (15) that

$$J_{SM}(\theta) \quad = \quad J_{DSM}(\theta) + C_1 - C_2,$$

which means that the optimization of J_{SM} and J_{DSM} are equivalent. □

Some references on Schrödinger bridge and applications to generative modeling

A. Alouadi, B. Barreau, L. Carlier, H. Pham : Robust time series generation via Schrödinger bridge : a comprehensive evaluation, 2025, arXiv : 2503.02943

Y. Chen, T. Georgiou, M. Pavon : Stochastic control liaisons : Richard Sinkhorn meets Gaspard Monge on a Schrödinger bridge, 2021, *SIAM Review, 63(2)*

C. Léonard : A survey of the Schrödinger problem and some of its connections with optimal transport, 2014, *Discrete Continuous and Dynamical systems, 34(4), 1533-1574*

G. Peyré, M. Cuturi : Computational optimal transport, 2019, *Foundations and trends in Machine learning*

V. De Bortoli, J. Thornton, J. Heng, A. Doucet : Diffusion Schrödinger bridge with applications to score-based generative modeling, 2021, *NeurIPS*

G. Wang, Y. Jiao, Q. Xu, Y. Wang, C. Yang : Deep generative learning via Schrödinger bridge, 2021, *ICML*

M. Hamdouche, P. Henry Labordère, H. Pham : Generative modeling for time series via Schrödinger bridge, 2023, arXiv :2304.06093